

## ALL-SKY INTERFEROMETRY WITH SPHERICAL HARMONIC TRANSIT TELESCOPES

J. RICHARD SHAW<sup>1,†</sup>, KRIS SIGURDSON<sup>2</sup>, UE-LI PEN<sup>1</sup>, ALBERT STEBBINS<sup>3</sup>, AND MICHAEL SITWELL<sup>2</sup><sup>1</sup>Canadian Institute for Theoretical Astrophysics, 60 St. George St., Toronto, ON M5S 3H8, Canada<sup>2</sup>Department of Physics and Astronomy, University of British Columbia, Vancouver, BC V6T 1Z1, Canada<sup>3</sup>Theoretical Astrophysics Group, Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*Draft version February 5, 2013*

## ABSTRACT

In this paper we describe the spherical harmonic transit telescope, a novel formalism for the analysis of transit radio telescopes. This all-sky approach bypasses the curved sky complications of traditional interferometry and so is particularly well suited to the analysis of wide-field radio interferometers. It enables compact and computationally efficient representations of the data and its statistics that allow new ways of approaching important problems like mapmaking and foreground removal. In particular, we show how it enables the use of the Karhunen-Loève transform as a highly effective foreground filter, suppressing realistic foreground residuals for our fiducial example by at least a factor twenty below the 21 cm signal even in highly contaminated regions of the sky. This is despite the presence of the angle-frequency mode mixing inherent in real-world instruments with frequency-dependent beams. We show, using Fisher forecasting, that foreground cleaning has little effect on power spectrum constraints compared to hypothetical foreground-free measurements. Beyond providing a natural real-world data analysis framework for 21 cm telescopes now under construction and future experiments, this formalism allows accurate power spectrum forecasts to be made that include the interplay of design constraints and realistic experimental systematics with twenty-first century 21 cm science.

## 1. INTRODUCTION

Mapping the Universe with the 21 cm line of neutral hydrogen will revolutionise our view of the Universe. It holds the promise of unravelling the mysteries of dark energy (Chang et al. 2008; Loeb & Wyithe 2008), unveiling the epoch of reionisation (EoR) (Furlanetto et al. 2006), and perhaps even extending our view of the cosmos out far enough to shine light on the primordial dark ages (Loeb & Zaldarriaga 2004). Rapidly probing large volumes of the Universe requires new large wide-field telescopes along with powerful new digital processing hardware such as GMRT<sup>1</sup>, LOFAR<sup>2</sup>, MWA<sup>3</sup>, Omniscopes, PAPER<sup>4</sup>, BAOBAB<sup>5</sup>, BAORadio<sup>6</sup>, BINGO<sup>7</sup>, CHIME<sup>8</sup>, EMBRACE/EMMA<sup>9</sup> and Tianlai<sup>10</sup>. In recent years it has become increasingly clear that new methods of interpreting and analysing the data from these revolutionary new instruments will be necessary to realize their scientific potential (Myers et al. 2003; Tegmark & Zaldarriaga 2009; Parsons & Backer 2009; Liu et al. 2010; Liu & Tegmark 2011; Parsons et al. 2012; Dillon et al. 2012).

We describe here the spherical harmonic transit telescope, a new paradigm for analysing wide-field transit telescopes in the spherical harmonic domain that is naturally suited to mapping the 21 cm Universe. Any

telescope with fixed pointing observes the sky transit through its field of view. The rotation about the poles periodically over the course of a sidereal day creates a linear correspondence between time  $t$  and azimuthal angle  $\phi$ . We obtain a simple mapping between the observed data and a linear combination of the spherical harmonic coefficients  $a_{lm}$  of the sky at fixed angular wavenumber  $m$ , mediated by the angular response of each element. In what follows, we elaborate and make precise this basic idea in the context of wide-field interferometers, including the radial (frequency) direction.

This formalism diverges sharply with traditional characterizations of radio interferometry that are better suited to observations with a narrow field of view, often assuming tracking of a particular source of interest, and exploit the Fourier transform mapping between the sky and the  $uv$ -plane. What we describe here is an all sky formalism for describing interferometry that naturally incorporates the observable modes on the sky — the spherical harmonics.

The foremost challenge for any 21 cm mapping experiment is separating the cosmological signal from astrophysical contaminants which are  $10^3$ – $10^5$  times larger (Furlanetto et al. 2006; Morales & Wyithe 2010). Conceptually this is simple — the primary foreground sources (diffuse synchrotron emission from the Galaxy and emission from extragalactic point sources) are smooth as a function of frequency, while the 21 cm signal decorrelates quickly as each frequency corresponds to a different radial slice of the Universe. To remove foregrounds one just needs to model and remove the smooth frequency component from their observations. Unfortunately, in practice, the large dynamic range between the amplitude of the foregrounds and the 21 cm signal makes several real-world effects extremely problematic. While the properties of the cosmological 21 cm signal are thought

<sup>†</sup> jrs65@cita.utoronto.ca<sup>1</sup> <http://gmrt.ncra.tifr.res.in/><sup>2</sup> <http://www.lofar.org/><sup>3</sup> <http://www.mwatelescope.org/><sup>4</sup> <http://eor.berkeley.edu/><sup>5</sup> <http://bao.berkeley.edu/><sup>6</sup> <http://groups.lal.in2p3.fr/bao21cm/><sup>7</sup> Battye et al. (2012)<sup>8</sup> <http://chime.phas.ubc.ca/><sup>9</sup> Kant et al. (2011)<sup>10</sup> <http://tianlai.bao.ac.cn/>

to be well understood, the astrophysical foregrounds are poorly constrained at the small angular and frequency scales that will be probed by forthcoming 21 cm experiments — a successful technique should be robust to uncertainties in foreground modelling. Of course, these experiments will themselves help characterize the properties of real-world foregrounds. More troublesome is the phenomena of angular-frequency mode mixing: in any real experiment the shape of the beam on the sky will vary with the observed frequency (Liu et al. 2009). This mode mixing makes simple frequency only foreground removal methods ineffective in practice. We show below that the spherical harmonic transit telescope formalism can naturally address the issues of model uncertainty and mode mixing, and enable efficient and effective discrimination of the 21 cm signal from obscuring foregrounds.

Any foreground removal method aims to find a subset of the data within which there is significantly more 21 cm signal than astrophysical foregrounds. However, in the presence of mode-mixing, it is not obvious how to select a basis which separates the two components — what we would like is a method which can automatically generate it. Just such a technique exists in the form of the Karhunen-Loève (KL) transform. In this paper we show how the  $m$ -mode formalism, described hence, makes the use of the KL transform computationally feasible. The result is a remarkably effective and robust filter for rejecting bright foregrounds and we demonstrate its effectiveness using realistic simulations of the radio sky and a simple fiducial interferometer configuration.

In Sec. 2 we introduce the all-sky formalism that is the basis for this technique. In Sec 3 we discuss the map-making process in the spherical harmonic transit telescope paradigm. In Sec. 4 we discuss how to best represent statistics of the cosmological 21 cm signal and foregrounds in the measurement basis, and in Sec. 5 we discuss how the Karhunen-Loève transform can be used to detect faint signals in the presence of bright foregrounds. In Sec. 6 we quantify the information lost due to foreground removal using Fisher Analysis, and estimate errors on power spectra. We conclude in Sec. 7. In Appendix A we discuss the signal and foreground models we employ. In Appendix B we describe how we create realistic simulations of radio emission.

## 2. FORMALISM

In this section we introduce the  $m$ -mode formalism, a new description of the measurement process for transit interferometers.

In radio interferometry a visibility  $V_{ij}$  is the instantaneous correlation between two feeds  $F_i$  and  $F_j$ . We will assume that we can take a linear combination of the signal from a dual polarisation antenna, with no cross-polarisation or polarisation leakage, such that we are sensitive only to the total intensity (Stokes  $I$ ) part of the sky. The fully polarised extension to this work is also a tractable problem, we address this in a subsequent paper, Shaw et al. (2013). At any instant, a visibility is given by

$$\begin{aligned} V_{ij} &= \langle F_i F_j^* \rangle \\ &= \frac{1}{\Omega_{ij}} \int d^2 \hat{\mathbf{n}} A_i(\hat{\mathbf{n}}) A_j^*(\hat{\mathbf{n}}) e^{2\pi i \hat{\mathbf{n}} \cdot \mathbf{u}_{ij}} T(\hat{\mathbf{n}}) \end{aligned} \quad (1)$$

where  $\mathbf{u}_{ij} = (\mathbf{r}_i - \mathbf{r}_j)/\lambda$  is the spatial separation between the two feeds divided by the observed wavelength (that is the separation in the  $uv$ -plane),  $\hat{\mathbf{n}}$  is the position on the celestial sphere, and  $A_i(\hat{\mathbf{n}})$  gives the primary beam of feed  $i$ . In the above we have normalised our visibilities such that they are temperature like, and we have defined them in terms of the brightness temperature  $T = \lambda^2 I / 2k_b$  instead of the total intensity  $I$ . The quantity  $\Omega_{ij} = \sqrt{\Omega_i \Omega_j}$  is the geometric mean of the individual beam solid angles

$$\Omega_i = \int |A_i(\hat{\mathbf{n}})|^2 d^2 \hat{\mathbf{n}} \quad (2)$$

which also gives the effective antenna area  $A_{\text{eff}} \Omega = \lambda^2$ . This ensures that for a sky with uniform brightness temperature  $T$  the auto-correlation of an antenna  $V_{ii} = T$  with our definition.

As the Earth turns both the primary beams and the baseline separations rotate relative to the celestial sphere. This means the measured visibilities change periodically with the sidereal day. We take this into account by explicitly including the dependence on the azimuthal angle  $\phi$  and by averaging over each sidereal day.

The measured visibilities are also corrupted by instrumental noise for which we add a noise term  $n_{ij}(\phi)$ . We assume the noise is stationary such that its statistics are independent of  $\phi$ . Rewriting Equation (1) in terms of a transfer function  $B_{ij}$  leaves the measured visibility as

$$V_{ij}(\phi) = \int d^2 \hat{\mathbf{n}} B_{ij}(\hat{\mathbf{n}}; \phi) T(\hat{\mathbf{n}}) + n_{ij}(\phi) \quad (3)$$

where the transfer function is

$$B_{ij}(\hat{\mathbf{n}}; \phi) = \frac{1}{\Omega_{ij}} A_i(\hat{\mathbf{n}}; \phi) A_j^*(\hat{\mathbf{n}}; \phi) e^{2\pi i \hat{\mathbf{n}} \cdot \mathbf{u}_{ij}(\phi)}. \quad (4)$$

Taking advantage of the periodicity in  $\phi$ , we Fourier transform the system

$$V_m^{ij} = \int \frac{d\phi}{2\pi} V_{ij}(\phi) e^{-im\phi} \quad (5)$$

$$= \sum_{lm'} \int \frac{d\phi}{2\pi} B_{lm'}^{ij}(\phi) a_{lm'} e^{-im\phi} + n_m^{ij} \quad (6)$$

where to proceed to the second line we have inserted the spherical harmonic expansions of both the sky, and the beam transfer function

$$T(\hat{\mathbf{n}}) = \sum_{lm} a_{lm} Y_{lm}(\hat{\mathbf{n}}), \quad (7)$$

$$B_{ij}(\hat{\mathbf{n}}; \phi) = \sum_{lm} B_{lm}^{ij}(\phi) Y_{lm}^*(\hat{\mathbf{n}}). \quad (8)$$

Note that we have defined  $B_{lm}^{ij}$  relative to the conjugate spherical harmonic in order to simplify later notation. As the  $\phi$  dependence simply rotates the functions about the Earth's polar axis, the transfer function at any  $\phi$  is trivially  $B_{lm}^{ij}(\phi) = B_{lm}^{ij}(\phi=0) e^{im\phi}$ . Combined with the exponential factor in the integral, this simply generates the Kronecker delta  $\delta_{mm'}$ , and we find

$$V_m^{ij} = \sum_l B_{lm}^{ij} a_{lm} + n_m^{ij}. \quad (9)$$

This gives a simple description of how the observed sky maps into the measured data given a telescope design (which is contained in the beam transfer matrices  $B_{lm}^{ij}$ ). This transformation does not mix  $m$ -modes on the sky, and can therefore be performed on an  $m$ -by- $m$  basis — for any particular  $m$  and frequency  $\nu$  the measured visibilities are simply a projection of the  $l$ -modes on the sky for the measured  $m$ . As the optical system is of a finite size, this limits both the  $l$  and  $m$  to which the telescope is sensitive, ensuring we only need to consider a finite number of degrees of freedom, both measured ( $V_m$ ) and on the sky ( $a_{lm}$ ).

In fact whilst the positive and negative  $m$ -modes may be independent measurements they are still observations of the same sky — by transforming the conjugate  $V_{-m}^*$  and using that  $a_{lm} = a_{l,-m}^*$  for a real field we see that

$$V_{-m}^{ij*} = \sum_l (-1)^m B_{l,-m}^{ij*} a_{lm} + n_{-m}^{ij*}. \quad (10)$$

In light of this we will change our notation such that we are considering only the actual degrees of freedom on the sky. Let us separate out the positive and negative  $m$  parts by defining

$$B_{lm}^{ij,+} = B_{lm}^{ij} \quad n_m^{ij,+} = n_m^{ij} \quad (11)$$

$$B_{lm}^{ij,-} = (-1)^m B_{l,-m}^{ij*} \quad n_m^{ij,-} = n_{-m}^{ij*} \quad (12)$$

which is valid for  $m \geq 0$ . Additionally to prevent double counting the  $m = 0$  measurement we need to set  $B_{l0}^{ij-} = n_0^{ij-} = 0$ . This gives a modified version of Equation (14)

$$V_m^{ij,\pm} = \sum_l B_{lm}^{ij,\pm} a_{lm} + n_m^{ij,\pm}. \quad (13)$$

For brevity of notation, we will introduce a label  $\alpha$  which indexes both the positive and negative  $m$  parts of all included feed pairs  $ij$ , such that any particular  $\alpha$  specifies exactly the values of  $ij, \pm$  (exactly how  $\alpha$  is packed is unimportant). This gives

$$V_m^\alpha = \sum_l B_{lm}^\alpha a_{lm} + n_m^\alpha. \quad (14)$$

The beam transfer matrices above can be written in an explicit matrix notation

$$(\mathbf{B}_m)_{(\alpha\nu)(l\nu')} = B_{lm}^{\alpha,\nu} \delta_{\nu\nu'} \quad (15)$$

where the row index labels all combinations of baseline ( $\alpha$ ) and frequency ( $\nu$ ), whereas the column index is over all multipole ( $l$ ) and frequencies ( $\nu'$ ). Similarly we can define vectors for the visibilities and harmonic coefficients

$$(\mathbf{v}_m)_{(\alpha\nu)} = V_m^{\alpha,\nu} \quad (\mathbf{a}_m)_{(l\nu)} = a_{lm}^\nu. \quad (16)$$

From here onwards we'll drop the subscript  $m$  denoting the spherical harmonic order, all the equations below are valid for any  $m$ . This allows us to rewrite Equation (14) as

$$\mathbf{v} = \mathbf{B} \mathbf{a} + \mathbf{n}. \quad (17)$$

This simple linear description of the measurement process of a transit telescope is extremely powerful. By reducing it down to a linear mapping between a finite

number of degrees of freedom it allows us to apply the standard tools of signal processing. In the subsequent sections we apply it to solve two challenging problems in 21 cm radio astronomy.

### 3. MAP-MAKING

In astronomy being able to transform our measured signal into an accurate map of the sky is essential. Whilst in this paper we explicitly avoid this process for our analysis, preferring to carry it out directly in the data space, maps are still needed for visualisation and cross-checking. Map-making with interferometric data is generally a complicated process performed by algorithms such as CLEAN (Högbom 1974) and its derivatives. This is especially true with wide fields of view where mosaicing and  $w$ -projection are generally required. However, the  $m$ -mode formalism makes the map-making process on the full sky conceptually simple.

First, we assume that the instrumental noise  $\mathbf{n}$  follows a complex gaussian distribution with covariance  $\mathbf{N} = \langle \mathbf{n} \mathbf{n}^\dagger \rangle$ , and the different frequency channels are independent. For stationary noise, the  $m$ -modes are uncorrelated and the likelihood function of the observed sky for a single  $m$  and frequency  $\nu$  is

$$p(\mathbf{v}|\mathbf{a}) = \frac{1}{|\pi \mathbf{N}|} \exp \left( -(\mathbf{v} - \mathbf{B} \mathbf{a})^\dagger \mathbf{N}^{-1} (\mathbf{v} - \mathbf{B} \mathbf{a}) \right) \quad (18)$$

where the vector  $\mathbf{a}$  contains all harmonic coefficients for the given  $m$ .

To estimate the sky corresponding to a given set of visibilities we will look for a maximum likelihood solution  $dp/d\mathbf{a} = 0$ . In particular we want to find the value of  $\mathbf{a}$  that minimises

$$\chi^2 = \left| \mathbf{N}^{-\frac{1}{2}} \mathbf{v} - \left( \mathbf{N}^{-\frac{1}{2}} \mathbf{B} \right) \mathbf{a} \right|^2. \quad (19)$$

The matrix  $\mathbf{N}^{-\frac{1}{2}}$  represents any factorisation such that  $(\mathbf{N}^{-\frac{1}{2}})^\dagger \mathbf{N}^{-\frac{1}{2}} = \mathbf{N}^{-1}$ . Provided  $\mathbf{N}$  contains no noiseless modes, it is positive-definite and so this factorisation should exist. The maximum likelihood solution is given by the Moore-Penrose pseudo-inverse<sup>11</sup>

$$\hat{\mathbf{a}} = \left( \mathbf{N}^{-\frac{1}{2}} \mathbf{B} \right)^+ \mathbf{N}^{-\frac{1}{2}} \mathbf{v}, \quad (20)$$

where the superscript  $+$  denotes the pseudo-inverse.

Depending on the number of baselines measured and the maximum  $l$  we are sensitive to, the problem may either be over- or under-constrained. In either regime the Moore-Penrose pseudo-inverse gives a solution, in the former case this reduces to the standard map making equation  $\hat{\mathbf{a}} = (\mathbf{B}^\dagger \mathbf{N}^{-1} \mathbf{B})^{-1} \mathbf{B}^\dagger \mathbf{N}^{-1} \mathbf{v}$ , and in the latter case selects the solution which also minimises  $|\hat{\mathbf{a}}|^2$ , effectively setting unconstrained degrees of freedom to zero.

As both distinct frequencies and  $m$ -modes are independent, map-making for a set of full sky observations is a case of collating the estimates for each individual  $\nu$  and  $m$ .

<sup>11</sup> For details see [http://en.wikipedia.org/wiki/Moore-Penrose\\_pseudoinverse](http://en.wikipedia.org/wiki/Moore-Penrose_pseudoinverse)

#### 4. TWO POINT STATISTICS

For Intensity Mapping experiments, our data has three components: the 21 cm signal which we are ultimately trying to extract, the foregrounds, and instrumental noise. Understanding the 2-point statistics of the data is of paramount importance to our analysis — not only do the correlations of the signal encode most of the cosmological information that we are interested in (see Appendix A), but to efficiently extract this we require knowledge of the 2-point statistics of all three components. Here we write down the linear relationship between these 2-point statistics of the data, and how they are related to the underlying physical correlations.

The statistics of instrumental noise live in the visibility space, the basis of our measurements. However the other components are naturally represented on the sky, and must be projected into this space using Equation (14). The lowest non-zero moment of the visibilities is their covariance

$$\begin{aligned} C_{(\alpha\nu m);(\alpha'\nu'm')} &= \langle V_{\alpha\nu}^m V_{\alpha'\nu'}^{m*} \rangle \\ &= \sum_{ll'} B_{lm}^{\alpha\nu} \langle a_{lm\nu}^* a_{l'm'\nu'} \rangle B_{l'm'}^{\alpha'\nu'*} + \langle n_{(\alpha\nu m)} n_{(\alpha'\nu'm')}^* \rangle. \end{aligned} \quad (21)$$

This is the covariance between all measured degrees of freedom: baselines, frequencies, and  $m$ -modes. For the experiments listed in Section 1 we expect  $\gtrsim 10^3$ ,  $\sim 10^2$  and  $10^3$  respectively. This gives matrices of dimension  $\gtrsim 10^8$ , too large to be tackled with current technology, both in terms of computation and storage.

Instead, let us make an approximation that will dramatically reduce this complexity. If we think of the sky as a statistically isotropic random field, its two point statistics become dramatically simpler

$$\langle a_{lm\nu} a_{l'm'\nu'}^* \rangle = C_l(\nu, \nu') \delta_{ll'} \delta_{mm'}, \quad (22)$$

and importantly, they are automatically uncorrelated in the  $m$  index. This means that the full signal covariance Equation (21) is block diagonal and thus allows us to calculate all statistics on an  $m$ -by- $m$  basis. For a specific  $m$ -mode

$$C_{(\alpha\nu);(\alpha'\nu')} = \sum_l B_l^{\alpha\nu} B_l^{\alpha'\nu'*} C_l(\nu, \nu') + N_{(\alpha\nu);(\alpha'\nu')} , \quad (23)$$

where we have dropped all the  $m$ -indices, and  $N$  is the power spectrum on the instrumental noise. As the number of  $m$ -modes we are sensitive to is usually  $\gtrsim 10^3$ , assuming statistical isotropy saves at least a factor of a million in computation and a thousand in storage. This ability to efficiently perform calculations incorporating the full statistics opens up new avenues for the data analysis of transit instruments. Synchrotron emission from our galaxy clearly violates this assumption of statistical isotropy, though, as we will demonstrate, this does not appear to affect our analysis and in particular our ability to clean foregrounds.

In matrix notation

$$\mathbf{C} = \mathbf{B} \mathbf{C}_{\text{sky}} \mathbf{B}^\dagger + \mathbf{N} . \quad (24)$$

where we will split  $\mathbf{C}_{\text{sky}}$  into independent 21 cm signal and foreground parts  $\mathbf{C}_{\text{sky}} = \mathbf{C}_{21} + \mathbf{C}_f$ .

The statistical models used for each component are chosen to be appropriate for the frequency ranges of interest. In the fiducial example that follows this is 400–600 MHz, corresponding to  $z \sim 1$ –2 for the cosmological signal. These are described in Appendix A.

#### 5. FOREGROUND REMOVAL WITH THE KARHUNEN-LOÈVE TRANSFORM

To clean our data we simply aim to find a subset within which there is significantly more 21 cm signal than the astrophysical foregrounds. However, in the presence of mode mixing there is no immediately apparent representation in which to perform this. This basis can be found using the Karhunen-Loève transform (often called the Signal-to-Noise eigendecomposition), which has a long history in Cosmology (e.g. Bond 1995; Vogeley & Szalay 1996) and has been used for the analogous problem of E/B mode separation for polarisation of the CMB (Lewis et al. 2002; Bunn et al. 2003). This transform simultaneously diagonalises both the signal and foreground covariance matrices, generating a set of modes with no foreground or signal correlations. This makes comparing the amount of signal and foreground power in each mode trivial.

To reduce the risk of foreground uncertainties biasing our analysis, we will prioritise the removal of foreground contaminated modes at the expense of cosmological signal. In contrast, the instrumental noise is well understood, and we should be able to dig deeper into this contaminant to extract useful cosmological information with little risk. In practice, this means we will start with a filter which aggressively removes foregrounds only; subsequently we will add back in the instrumental noise, which will allow us to compress the data by removing completely noise-dominated modes, while retaining those with a small fraction of signal.

This requires models for the statistics of both the signal and the foregrounds. The signal is modelled as a simple gaussian random field for the 21 cm emission, whereas the foreground model includes both the synchrotron emission from our galaxy, and the contribution from a background of extragalactic point sources. The details of both are discussed in Appendix A.

The Karhunen-Loève transform seeks to find a linear transformation of the data  $\mathbf{v}' = \mathbf{P}\mathbf{v}$  such that the 21 cm signal  $\mathbf{S} = \mathbf{B} \mathbf{C}_{21} \mathbf{B}^\dagger$  and foreground  $\mathbf{F} = \mathbf{B} \mathbf{C}_f \mathbf{B}^\dagger$  covariance matrices are jointly diagonalised. That is

$$\mathbf{S} \rightarrow \mathbf{S}' = \mathbf{P} \mathbf{S} \mathbf{P}^\dagger = \mathbf{\Lambda} , \quad (25)$$

and

$$\mathbf{F} \rightarrow \mathbf{F}' = \mathbf{P} \mathbf{F} \mathbf{P}^\dagger = \mathbf{I} , \quad (26)$$

where  $\mathbf{\Lambda}$  is a diagonal matrix, and  $\mathbf{I}$  is the identity. In this diagonal basis we can simply compare the amount of power expected in each mode by the ratio of the diagonal elements (this is given by the corresponding entries of  $\mathbf{\Lambda}$ ), and identify the regions of the space with low foreground contamination (large entries in  $\mathbf{\Lambda}$ ).

This transformation can be found by solving the generalised eigenvalue problem  $\mathbf{S}\mathbf{x} = \lambda\mathbf{F}\mathbf{x}$ . This gives a set of eigenvectors  $\mathbf{x}$ , and corresponding eigenvalues  $\lambda$ . Writing the eigenvectors in a matrix  $\mathbf{P}$ , row-wise, gives the transformation matrix to diagonalise the covariances.

The eigenvalues  $\lambda$  corresponding to each eigenvector give the diagonal matrix  $\Lambda$ .

To isolate the 21 cm signal, we want select modes with eigenvalue (signal-to-foreground power) greater than some threshold (see Figure 1). To project into this basis we define the matrix  $\mathbf{P}_s$  which contains only the rows from  $\mathbf{P}$  corresponding to eigenvalues greater than the threshold  $s$ .

For most analysis we can work directly in the eigenbasis. However, for visualising our results, we want to be able to transform back to the sky (by way of the measured visibilities). To project back into the higher dimensional space we simply generate the full inverse  $\mathbf{P}^{-1}$  and remove columns corresponding to the rejected modes (we denote this matrix  $\tilde{\mathbf{P}}_s$ ). This is equivalent to projecting into the full eigenbasis, zeroing the foreground contaminated modes, and then using the full-inverse  $\mathbf{P}^{-1}$ .

For further analysis, we must include all noise terms, both foregrounds and instrumental. Writing the total noise contribution as  $\mathbf{N}_{\text{all}} = \mathbf{F} + \mathbf{N}$ , the matrix in the truncated basis is

$$\mathbf{N}^{\text{all}} \rightarrow \mathbf{N}_s^{\text{all}} = \mathbf{P}_s (\mathbf{F} + \mathbf{N}) \mathbf{P}_s^\dagger \quad (27)$$

$$= \mathbf{I} + \mathbf{P}_s \mathbf{N} \mathbf{P}_s^\dagger. \quad (28)$$

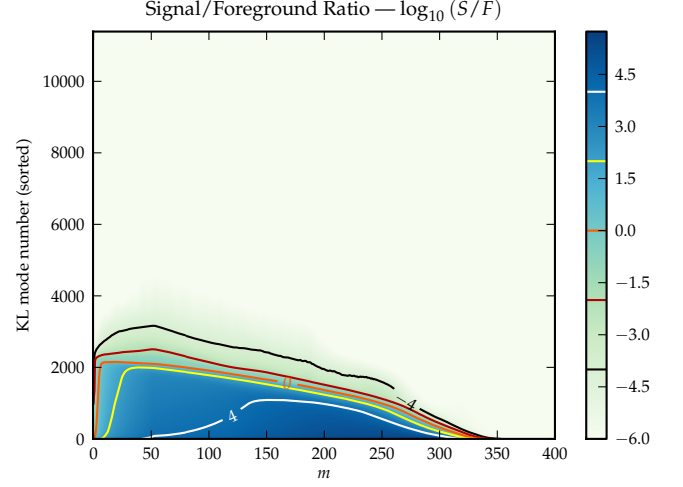
As the transformed instrumental noise matrix will not remain diagonal this gives a correlated component between all our modes. However, as it is useful if our modes are uncorrelated we make a further KL-transformation on the foreground removed signal  $\mathbf{S}_s = \Lambda_s$ , and total noise  $\mathbf{N}_s^{\text{all}}$  covariance matrices. For computational and storage efficiency we apply a further cut-off to include only modes with a signal to total-noise ratio greater than cut-off value  $t$ . We denote this projection matrix  $\tilde{\mathbf{Q}}$ . For notational convenience we will write the total transformation in terms of a single matrix  $\mathbf{R} = \mathbf{Q}_t \mathbf{P}_s$ , having chosen suitable values for the two cut-offs  $s$  and  $t$ . As above, we will define an inverse  $\tilde{\mathbf{R}} = \tilde{\mathbf{P}}_s \tilde{\mathbf{Q}}_t$  which remains orthogonal to the removed space. Quantities in this final basis we denote with tildes, for example a visibility mapped into this basis is  $\tilde{\mathbf{v}} = \mathbf{R} \mathbf{v}$ , and a covariance is  $\tilde{\mathbf{C}} = \mathbf{R} \mathbf{C} \mathbf{R}^\dagger$ .

### 5.1. Cylinder Example

While this method works with any transit telescope, to illustrate the foreground removal process we will simulate a cylinder telescope, such as CHIME or the Pittsburgh Cylinder Telescope (Bandura 2011). These are transit interferometers composed of multiple parabolic cylinders where each only focuses in the East-West direction. This gives a long and thin primary beam on the sky, extending nearly from horizon to horizon in the North-South direction but which is only around 1 degree wide East-West.

Feeds are spaced along the axis of each cylinder — when correlated these provide resolution in the N-S direction. Correlations between cylinders enhance the E-W resolution. The telescope operates as a transit telescope such that the entire visible sky is observed once per sidereal day.

For a cylinder uniformly illuminated by a particular feed, near the axis the beam pattern is a sinc function in the E-W direction, and uniform in the N-S direction



**Figure 1.** The Signal-to-Foreground spectrum for all  $m$ -modes. We have plotted  $\log_{10} \lambda_{im}$ , where for each  $m$  the eigenvalues have been sorted by in ascending order (thus there is no physical interpretation to the vertical direction). The contours are drawn at  $-4$ ,  $-2$ ,  $0$ ,  $2$  and  $4$ .

(Wilson et al. 2009, chapter 6). To extend this off-axis we modulate by projected area of the telescope giving

$$A^2(\hat{\mathbf{n}}) = \text{sinc}^2 \left( \pi \hat{\mathbf{n}} \cdot \hat{\mathbf{u}} \frac{W}{\lambda} \right) \Theta(\hat{\mathbf{n}} \cdot \hat{\mathbf{z}}) \hat{\mathbf{n}} \cdot \hat{\mathbf{z}} \quad (29)$$

where  $W$  is the cylinder width and  $\hat{\mathbf{z}}$  is a unit vector pointing to the zenith and  $\hat{\mathbf{u}}$  is a unit vector pointing East in the ground-plane. The step function  $\Theta$  masks out the regions where the sky is below the horizon, and the final factor  $\hat{\mathbf{n}} \cdot \hat{\mathbf{z}}$  accounts for the projected area of the telescope.

We model a two cylinder telescope observing the sky with 64 frequency channels from 400–600 MHz. Each cylinder is 15 m wide and has 60 feeds regularly spaced by 0.25 m (with the feeds lining up E-W between cylinders). The telescope is located at a latitude of  $45^\circ$ . These specifications correspond to a slightly smaller half-bandwidth version of the CHIME pathfinder telescope being constructed at DRAO.

The noise covariance is diagonal for all  $m$ , frequencies and baselines. For small  $m$ -modes with  $m \ll 1/(2\pi\Delta\phi)$  (where  $\Delta\phi$  is the angular integration time), the noise variance is

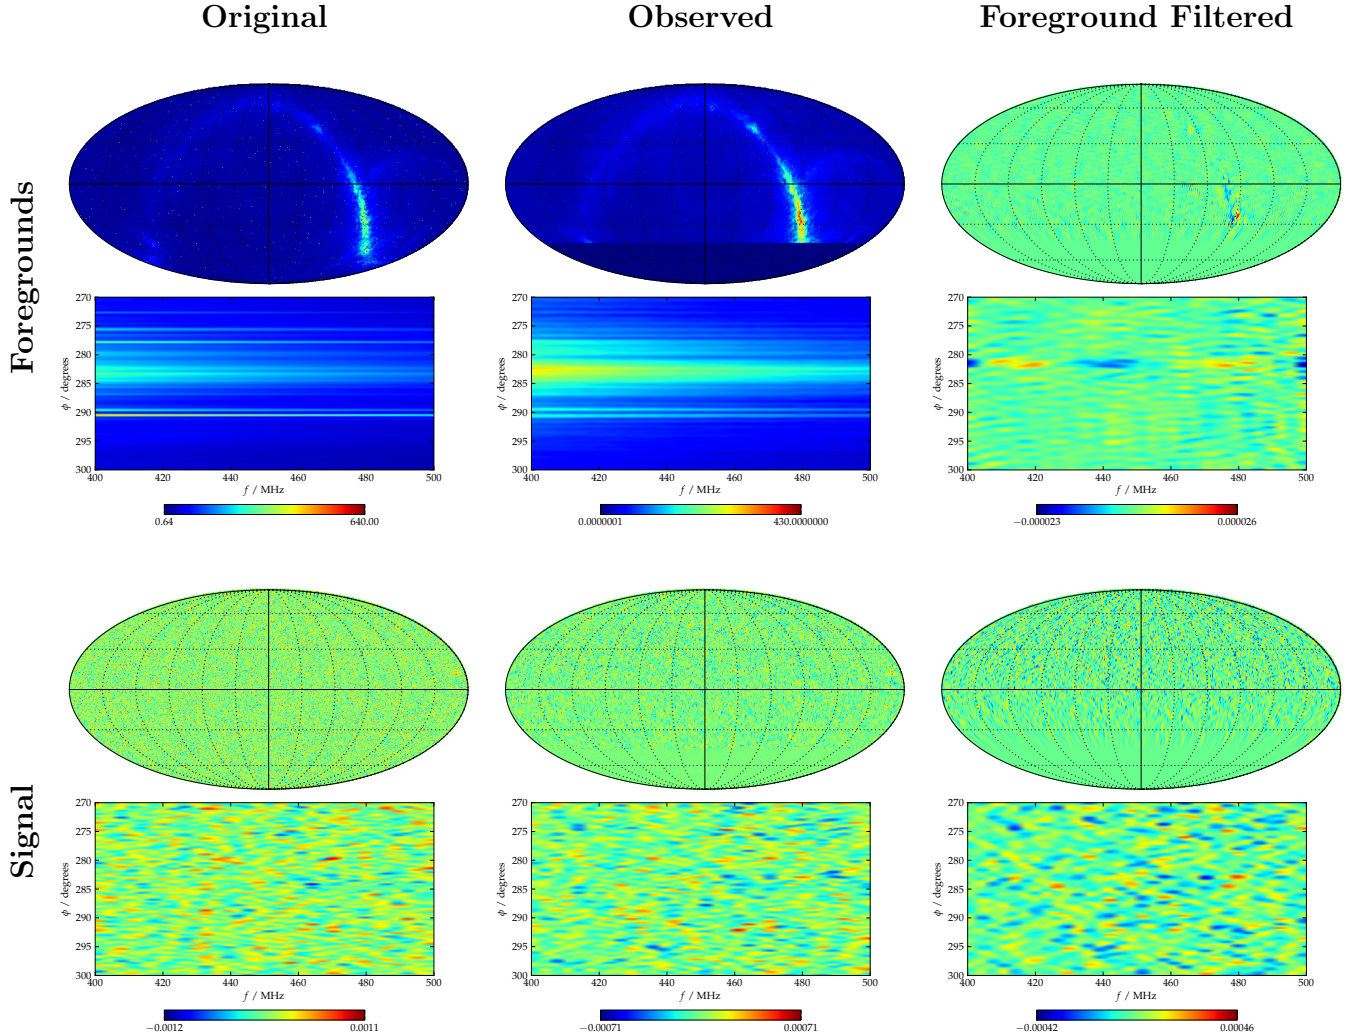
$$N_m^{ij} = \frac{T_{\text{sys},i}(\nu) T_{\text{sys},j}(\nu)}{4\pi N_{\text{day}} t_{\text{sid}} \Delta\nu}, \quad (30)$$

where  $T_{\text{sys},i}$  is the system temperature of a single polarisation of feed  $i$ ,  $N_{\text{day}}$  is the number of sidereal days observed,  $t_{\text{sid}}$  is the length of a sidereal day, and  $\Delta\nu$  is the width of the frequency channel. As we combine the two polarisations into a single unpolarised signal this reduces the noise power spectrum by a factor of two. For this example  $T_{\text{sys}} = 50$  K, and we assume two full years of observation (that is 730 complete sidereal days).

In Figure 1 we show the spectrum of Signal-to-Foreground eigenvalues for the telescope. The KL mode distribution of S/F has an extremely rapidly rising spectrum so that the information retained (approximately the number of modes) is rather insensitive to the cut threshold  $s$  for values between  $10^{-2}$ – $10^2$ .

To demonstrate the foreground removal process we





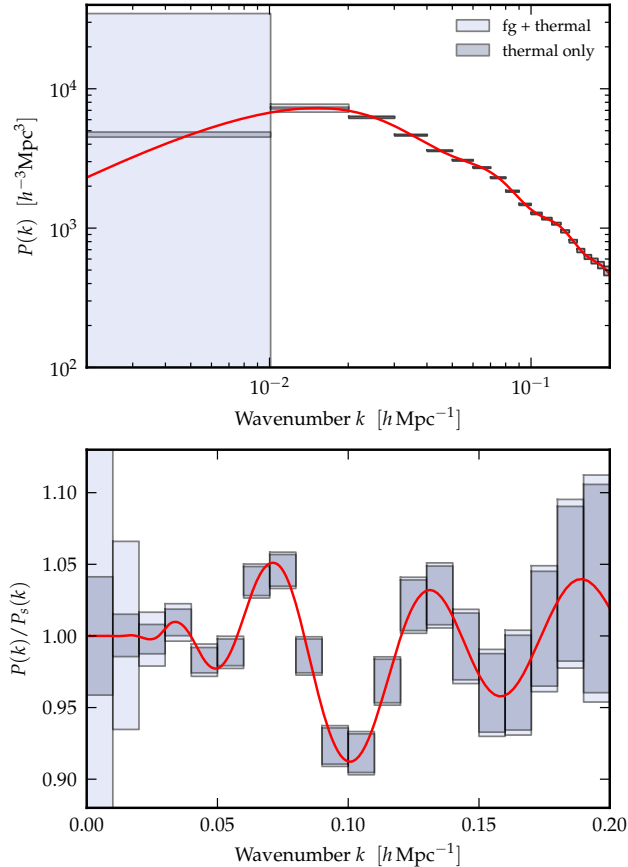
**Figure 2.** This plot illustrates the foreground removal process in action on simulations of the foregrounds-only (top row) and signal-only (bottom row). Each plot has two elements, an image of the 400 MHz frequency slice on top, and beneath, a cut through the celestial equator (from 270–300 degrees) showing the frequency axis. The left-most column shows the original simulations on the sky. The band appearing in the foreground frequency slice is the galactic plane. The middle column shows the maximum likelihood map that we would make from the measured visibilities without subtracting the low S/F modes. The maps are blank below  $\delta = -45^\circ$  because this area is always below the horizon for the telescope at a latitude of  $45^\circ$ . The final column shows the maps made after the foreground removal process (in this case we have discarded modes with  $S/F < 10$ ). This leaves a clear correspondence between the original signal simulation and the foreground subtracted signal, whilst leaving the foreground residuals over 20 times smaller in amplitude.

simulate time-streams from separate realisations of the signal and foregrounds using Equation (17), and project them through the filtering process to make maps. The visibilities are filtered using  $\mathbf{v}_{\text{clean}} = \mathbf{R}\mathbf{R}^\dagger\mathbf{v}$ , and then are turned into a 3D map using Equation (20). In Figure 2 we show the original simulation, the map made from the unfiltered visibilities, and the map made from the foreground filtered visibilities. The simulated signal and foreground maps are described in Appendix B. Note that the foreground maps are not simply realisations of the model used to generate the foreground filter — unlike the input model they are both non- gaussian and anisotropic. Figure 2 clearly illustrates how the foreground amplitude is dramatically reduced by the process, whilst the signal retains its overall character. Though the foreground residuals are clearly highest in the galactic centre, even these are significantly lower than the filtered signal.

## 6. FISHER ANALYSIS

In the previous section we have demonstrated that the Karhunen-Loève transform gives an effective method for removing foregrounds. Though a visual inspection of Figure 2 suggests that the 21 cm signal is largely untouched, we would like to be able to quantify how much useful information remains. In this section we will use the Fisher matrix (see Dodelson 2003, chap. 11 for an overview) to forecast power spectrum errors, for the same telescope, with and without foreground removal.

After projection into the reduced eigenbasis, let us assume that the remaining modes follow a complex gaussian distribution with zero mean. This assumption should be reasonable provided we have successfully removed the modes containing any significant foreground contribution. In this case the Fisher Information matrix



**Figure 3.** 21 cm Intensity Mapping provides a powerful technique for measuring the shape of the matter power spectrum. In the plots above we illustrate the power spectrum constraints that could be achieved with the large cylinder telescope. The top plot shows the constraints on the whole power spectrum, the lower plot zooms in on the region with the Baryon Acoustic Oscillations, dividing through by a smoothed spectrum to remove the general trend. The dark shaded bands are the errors we would find without foregrounds, where the only noise is instrumental, the light bands include both contributions.

for a set of parameters  $p_a$  is

$$F_{ab}^{(m)} = \text{tr} \left( \tilde{\mathbf{C}}_a \tilde{\mathbf{C}}^{-1} \tilde{\mathbf{C}}_b \tilde{\mathbf{C}}^{-1} \right). \quad (31)$$

where  $\tilde{\mathbf{C}}_a = \partial \tilde{\mathbf{C}} / \partial p_a$ . Though in the constructed eigenbasis  $\tilde{\mathbf{C}} = \tilde{\mathbf{\Lambda}} + \mathbf{I}$  is diagonal,  $\tilde{\mathbf{C}}_a$  can have off-diagonal elements. Again this process is performed on a per- $m$  basis. As there is no coupling between them, the total Fisher Information is simply the sum over all  $m$ -modes

$$F_{ab} = \sum_m F_{ab}^{(m)}. \quad (32)$$

For a set of parameters  $p_a$  that we are trying to determine, the inverse of the Fisher matrix is the lowest order approximation to their covariance.

In this work we will focus on forecasting the errors on the shape of the matter power spectrum  $P(k)$  whilst keeping all other cosmological parameters fixed. Such forecasting has been performed using the  $uv$ -plane in Seo et al. (2010) and Ansari et al. (2012).

We parametrise the power spectrum in terms of a linear

summation of different basis functions

$$P(k) = \sum_a p_a P_a(k). \quad (33)$$

In Appendix A we describe how to project this quantity into the angular power spectrum of 21 cm fluctuations that we use to calculate the visibility correlations. For simplicity, each of our bands is simply equal to the input power spectrum within a fixed  $k$ -band, and zero outside, such that the fiducial model is  $p_a = 1$ .

For the band-powers  $p_a$  that we are trying to estimate, the matrices  $\tilde{\mathbf{C}}_a$  are simply the projection of the basis functions  $P_a(\mathbf{k})$  into the eigenbasis. Starting from the angular power spectra  $C_{a;l}(\nu, \nu')$  corresponding to each of the basis functions  $P_a(\mathbf{k})$  (using Equation (A2))

$$\tilde{\mathbf{C}}_a = \mathbf{R} \mathbf{B} \mathbf{C}_{21,a} \mathbf{B}^\dagger \mathbf{R}^\dagger. \quad (34)$$

In practice explicitly calculating the  $\tilde{\mathbf{C}}_a$  this way is computationally expensive, we instead use a Monte-Carlo technique. We can form the estimator  $\hat{q}_a = \tilde{\mathbf{x}}^\dagger \tilde{\mathbf{C}}^{-1} \tilde{\mathbf{C}}_a \tilde{\mathbf{C}}^{-1} \tilde{\mathbf{x}}$ , which has the property that its covariance  $\langle \hat{q}_a \hat{q}_b \rangle - \langle \hat{q}_a \rangle \langle \hat{q}_b \rangle = F_{ab}$  (Padmanabhan et al. 2003). This means we can estimate the  $F_{ab}$  by averaging over realisations of  $\tilde{\mathbf{x}}$ . For details see Dillon et al. (2012).

In Figure 3 we plot the power spectrum errors for two cases: in the presence of foregrounds that have been cleaned using our method and without foregrounds at all. In the case without foregrounds,  $\mathbf{F} = 0$  and we only perform the final Karhunen-Loève transform to diagonalise the signal and instrumental noise. For the foregrounds we have cleaned modes with  $S/F < 10$  and additionally have removed modes with a small ratio of signal to total power. This corresponds to setting  $s = 10$  and  $t = 0.01$ . This is a clear demonstration of the effectiveness of the technique — it reduces our sensitivity on large scales as we would expect (as the removed foreground are smooth on large scales), while only slightly reducing our ability to constrain the small scale power spectrum.

## 7. CONCLUSION

In this paper we have introduced a powerful formalism for describing the measurement process of transit telescopes (either interferometric or otherwise). It is a natural formalism to describe interferometry on the full sky — sidestepping the standard complications that arise when dealing with wide field interferometric data such as mosaicing and  $w$ -projection. A spherical harmonic transit telescope allows for compact and computationally efficient representations of the data and its statistics, which enable new ways of approaching important problems like map-making and foreground removal.

Using the  $m$ -mode formalism and approximating the foregrounds as statistically isotropic allows the powerful Karhunen-Loève transformation to be used, automatically finding the basis in which the astrophysical foregrounds and 21 cm signal are maximally separated. The KL approach would be computationally impossible otherwise and is a key advantage of the  $m$ -mode formalism. Using this technique we can take the full three-dimensional dataset into account and overcome the mode-mixing problem. The filters we construct are highly effective and robust, a fact we have demonstrated

by propagating through realistically simulated 21 cm and foreground timestreams. In our fiducial example, shown in Figure 2, peak-to-peak foreground amplitude was reduced by a factor of  $\sim 2 \times 10^7$  leaving the peak-to-peak amplitude of the 21 cm signal around 20 times brighter than the foreground residuals.

We have also used this formalism to produce realistic forecasts for the power spectrum constraints from a fiducial 21 cm cylinder interferometer. We have demonstrated that foreground cleaning does not significantly degrade 21 cm power spectrum estimates on BAO scales and below compared to a hypothetical foreground-free measurement. We anticipate that the spherical harmonic transit telescope formalism will be a powerful tool that can be applied to inform experimental design and test the interplay between real-world systematics and design constraints on twenty-first century 21 cm science. We will explore this further in Shaw et al. (2013).

## ACKNOWLEDGEMENTS

We thank the CHIME team for stimulating discussions, and Matt Dobbs and Keith Vanderlinde for comments on an earlier version of this manuscript. KS, UP, and MS are supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada. The work of AS was supported by the DOE at Fermilab under Contract No. DE-AC02-07CH11359. KS thanks Perimeter Institute for Theoretical Physics for their hospitality. Computations were performed on the GPC supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto.

## APPENDIX

### A. SIGNAL AND FOREGROUND MODELS

We model the 21 cm signal and foregrounds as isotropic fields described by an angular power spectrum  $C_l(\nu, \nu')$ . We base our models on Santos et al. (2005), although we will only include the galactic synchrotron and extragalactic point source contributions. Both these contributions are assumed to take the form

$$C_l(\nu, \nu') = A \left( \frac{l}{100} \right)^{-\alpha} \left( \frac{\nu \nu'}{\nu_0^2} \right)^{-\beta} e^{-\frac{1}{2\xi_l^2} \ln^2(\nu/\nu')}. \quad (\text{A1})$$

As the models are calibrated for observations of the reionisation epoch, we need to transform them into the higher frequencies we are concerned with. We list the parameters for both these models in Table 1.

For the point source model, which is based on the results of Di Matteo et al. (2002), we change the pivot frequency  $\nu_0$  from 150 MHz to 408 MHz and also rescale the amplitude in order to raise the maximum flux of unsubtracted sources from 0.1 mJy to 0.1 Jy.

The galactic synchrotron model we use is not only calibrated for low frequencies but also high galactic latitudes. As we will measure large fractions of the sky we take this into account by changing the  $A$  and angular power-law

**Table 1**

Our model for the angular power spectrum  $C_l(\nu, \nu')$  is based on those of Santos et al. (2005) however we have adapted the parameters to better suit the full-sky intensity mapping regime we are interested in.

	$A \text{ (K}^2\text{)}$	$\alpha$	$\beta$	$\zeta$
Galaxy	$6.6 \times 10^{-3}$	2.80	2.8	4.0
Point Sources	$3.55 \times 10^{-4}$	2.10	1.1	1.0

index  $\beta$  to be consistent with the angular power spectrum of the 408 MHz Haslam map for galactic latitudes  $|b| > 5^\circ$  from La Porta et al. (2008).

We model the 21cm brightness temperature as being a biased tracer of the underlying matter fluctuations. These fluctuations are naturally characterised by the angular power spectrum (Lewis & Challinor 2007; Datta et al. 2007). However exact calculation of this quantity requires double-integration over highly oscillatory functions, instead we use the flat-sky approximation from Datta et al. (2007)

$$C_l(z, z') = \frac{1}{\pi \chi \chi'} \int_0^\infty dk_\parallel \cos(k_\parallel \Delta\chi) P_{T_b}(\mathbf{k}; z, z') \quad (\text{A2})$$

where  $\chi$  and  $\chi'$  are the comoving distances to redshift  $z$  and  $z'$ . Their difference is denoted by  $\Delta\chi = \chi - \chi'$ . The vector  $\mathbf{k}$  has the components  $k_\parallel$  and  $l/\bar{\chi}$  in the directions parallel and perpendicular to the line of sight ( $\bar{\chi}$  is the mean of  $\chi$  and  $\chi'$ ). This approximation is accurate to the 1% level for  $l > 10$  (Datta et al. 2007).

We model the 21cm brightness temperature power spectrum  $P_{T_b}$  as

$$P_{T_b}(\mathbf{k}; z, z') = \bar{T}_b(z) \bar{T}_b(z') (b + f\mu^2)^2 P_m(k; z, z') \quad (\text{A3})$$

where  $P_m(k; z, z') = P(k) D_+(z) D_+(z')$  is the real-space matter power spectrum,  $D_+(z)$  is the growth factor normalised such that  $D_+(0) = 1$ ,  $b$  is the bias, and the growth rate  $f = d \ln D_+ / d \ln a$ , the logarithmic derivative of the growth factor  $D_+$ . We assume that the bias  $b = 1$  at all redshifts. The mean brightness temperature is assumed to take the form

$$\bar{T}_b(z) = 0.3 \left( \frac{\Omega_{\text{HI}}}{10^{-3}} \right) \times \left( \frac{\Omega_m + (1+z)^{-3} \Omega_\Lambda}{0.29} \right)^{-1/2} \left( \frac{1+z}{2.5} \right)^{1/2} \text{ mK} \quad (\text{A4})$$

given in Chang et al. (2008). We assume that the neutral hydrogen fraction takes a value  $\Omega_{\text{HI}} = 5 \times 10^{-4}$  (Masui et al. 2013).

### B. SIMULATING ALL-SKY RADIO EMISSION

#### B.1. Galactic Synchrotron

In order to test our methods we require simulated maps of the Galactic emission from our own galaxy in the range 400–1400 MHz with 1 MHz resolution. Though there are maps at both 800 MHz and 1420 MHz, the only public all-sky radio survey in this range is the 408 MHz Haslam map (Haslam et al. 1982). However, the Global Sky Model (de Oliveira-Costa et al. 2008) is based on a compilation of maps from 10 MHz to 94 GHz. We use the



Global Sky Model to generate maps at both 400 MHz and 1420 MHz, and use these to estimate an effective spectral at each location on the sky

$$\alpha(\hat{n}) = \frac{\log T_{1420}(\hat{n}) - \log T_{400}(\hat{n})}{\log 1420 - \log 400}. \quad (\text{B1})$$

By combining this with the Haslam map<sup>12</sup> we can extrapolate to simulate a map of the sky at any desired set of frequencies

$$T_{\text{base}}(\hat{n}, \nu) = T_{408}(\hat{n}) \left( \frac{\nu}{408 \text{ MHz}} \right)^{\alpha(\hat{n})}. \quad (\text{B2})$$

Unfortunately this map lacks both the small scale angular fluctuations (because of the limited resolution of the Haslam map) and any spectral variations (because of the power law extrapolation) that would be present on the real sky. It is essential to include these to accurately test any foreground removal method.

To include these fluctuations we could add gaussian realisations of Equation (A1) (with the galactic synchrotron parameters, see Table 1) to the base map, which contain frequency and angular fluctuations at arbitrary resolutions. However the Haslam map already constrains what the sky looks like on scales  $\gtrsim 1^\circ$ , and the extrapolation with the spectral index map, is a constraint on the sky at 1420 MHz (on scales larger than  $5.1^\circ$ , the resolution of the Global Sky Model). Therefore we would like the combined simulated map to be consistent with these observations. We can do this by constraining the realisations to ensure there are no fluctuations on the scales constrained. In practice we do this by manipulating the amplitudes of the two highest valued eigenmodes of  $C_l(\nu, \nu')$  (from Equation A1) in each realisation, to ensure that the 408 MHz and 1420 MHz slices are zero when smoothed on  $1^\circ$  and  $5.1^\circ$  scales respectively.

A further problem is that we know the amplitude of small scale fluctuations varies over the sky, however our realisations are statistically isotropic. This is clearly demonstrated in the analysis of La Porta et al. (2008), which shows that the amplitude of the angular power spectrum traces the galactic structure. To reproduce this we use the RMS amplitude of fluctuations across the Haslam map in  $\sim 4^\circ$  patches (corresponding to Healpix pixels with NSIDE = 16), to rescale the fluctuations. In particular this generates an angular power spectrum on the whole sky which is consistent with a single power-law even when crossing through the beam-scale of the Haslam map into the simulated fluctuations. We do not include variations of the power-law index of the angular power spectrum as there appears to be no structure to the small variations found in La Porta et al. (2008).

### B.2. Extra-Galactic Point Sources

Our point source maps are constructed from two components, a population of bright point sources ( $S > 0.1$  Jy at 151 MHz) simulated directly, and a background of dimmer unresolved point sources ( $S < 0.1$  Jy) modelled as a gaussian random field.

The former is constructed directly by drawing from the point source distribution of Di Matteo et al. (2002),

each sourced is modelled as having pure power law emission with a random spectral index. The sources are distributed randomly over the sky. Very bright sources ( $S > 10$  Jy) are assumed to have been subtracted such that their residuals are less than this threshold.

The unresolved background is simulated by drawing a gaussian realisation from Equation (A1) with the point source model detailed in Table 1.

### B.3. 21 cm Signal

Simulations of the Cosmological 21cm emission are performed by drawing gaussian realisations from the flat-sky angular power spectrum (calculated using Equation (A2)).

## REFERENCES

- Ansari, R., Campagne, J. E., Colom, P., et al. 2012, A&A, 540, A129
- Bandura, K. 2011, PhD thesis, Carnegie Mellon University
- Battye, R. A., Browne, I. W. A., Dickinson, C., et al. 2012, ArXiv e-prints, arXiv:1209.0343 [astro-ph.CO]
- Bond, J. R. 1995, Physical Review Letters, 74, 4369
- Bunn, E. F., Zaldarriaga, M., Tegmark, M., & de Oliveira-Costa, A. 2003, Phys. Rev. D, 67, 023501
- Chang, T.-C., Pen, U.-L., Peterson, J. B., & McDonald, P. 2008, Physical Review Letters, 100, 091303
- Datta, K. K., Choudhury, T. R., & Bharadwaj, S. 2007, MNRAS, 378, 119
- de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. 2008, MNRAS, 388, 247
- Di Matteo, T., Perna, R., Abel, T., & Rees, M. J. 2002, ApJ, 564, 576
- Dillon, J. S., Liu, A., & Tegmark, M. 2012, ArXiv e-prints, arXiv:1211.2232 [astro-ph.CO]
- Dodelson, S. 2003, Modern cosmology
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, Phys. Rep., 433, 181
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, A&AS, 47, 1
- Högbom, J. A. 1974, A&AS, 15, 417
- Kant, G. W., Patel, P. D., Wijnholds, S. J., Ruiter, M., & van der Wal, E. 2011, IEEE Transactions on Antennas and Propagation, 59, 1990
- La Porta, L., Burigana, C., Reich, W., & Reich, P. 2008, A&A, 479, 641
- Lewis, A., & Challinor, A. 2007, Phys. Rev. D, 76, 083005
- Lewis, A., Challinor, A., & Turok, N. 2002, Phys. Rev. D, 65, 023505
- Liu, A., & Tegmark, M. 2011, Phys. Rev. D, 83, 103006
- Liu, A., Tegmark, M., Morrison, S., Lutomirski, A., & Zaldarriaga, M. 2010, MNRAS, 408, 1029
- Liu, A., Tegmark, M., & Zaldarriaga, M. 2009, MNRAS, 394, 1575
- Loeb, A., & Wyithe, J. S. B. 2008, Physical Review Letters, 100, 161301
- Loeb, A., & Zaldarriaga, M. 2004, Physical Review Letters, 92, 211301
- Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, ApJ, 763, L20
- Morales, M. F., & Wyithe, J. S. B. 2010, ARA&A, 48, 127
- Myers, S. T., Contaldi, C. R., Bond, J. R., et al. 2003, ApJ, 591, 575
- Padmanabhan, N., Seljak, U., & Pen, U. L. 2003, New A, 8, 581
- Parsons, A., Pober, J., McQuinn, M., Jacobs, D., & Aguirre, J. 2012, ApJ, 753, 81
- Parsons, A. R., & Backer, D. C. 2009, AJ, 138, 219
- Santos, M. G., Cooray, A., & Knox, L. 2005, ApJ, 625, 575
- Seo, H.-J., Dodelson, S., Marriner, J., et al. 2010, ApJ, 721, 164
- Shaw, J. R., Sigurdson, K., Stebbins, A., Pen, U.-L., & Sitwell, M. 2013, in preparation
- Tegmark, M., & Zaldarriaga, M. 2009, Phys. Rev. D, 79, 083530
- Vogele, M. S., & Szalay, A. S. 1996, ApJ, 465, 34
- Wilson, T. L., Rohlfs, K., & Hüttemeister, S. 2009, Tools of Radio Astronomy (Springer-Verlag)

<sup>12</sup> We use the map from the Legacy Archive for Microwave Background Data Analysis (LAMBDA), which has been processed to remove bright point sources and striping. See [http://lambda.gsfc.nasa.gov/product/foreground/haslam\\_408.cfm](http://lambda.gsfc.nasa.gov/product/foreground/haslam_408.cfm)